



NRL/FR/5523--96-9815

Real-time Packet Voice Traffic Support for Integrated Mobile Military Internetworking

JOSEPH P. MACKER

*Communications Systems Branch
Information Technology Division*

May 8, 1996

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE May 8, 1996		3. REPORT TYPE AND DATES COVERED
4. TITLE AND SUBTITLE Real-time Packet Voice Traffic Support for Integrated Mobile Military Internetworking			5. FUNDING NUMBERS PE - 0603792N 3995WXE8011	
6. AUTHOR(S) Joseph P. Macker				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory Washington, DC 20375-5320			8. PERFORMING ORGANIZATION REPORT NUMBER NRL/FR/5523--96-9815	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Space and Naval Warfare Systems Command Washington, Dc 20363-5100			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The Naval Research Laboratory (NRL) has designed and developed a flexible, low data rate network voice terminal application and a set of novel network integration techniques for supporting real-time packet voice over limited bandwidth wireless networks. The NRL Interactive Voice Exchange (IVOX) network voice application was developed to support low data rate (600-2400 bps) voice calls over a connectionless network. IVOX has been further modified to use a quality-of-service application programming interface (API) developed under the Data and Voice Integration Advanced Technology Demonstration (DVI ATD). Through use of the API, the network voice application can negotiate, establish, and maintain virtual circuit connection service for real-time, guaranteed delivery of network datagram traffic. Additional networking enhancements developed to support packet voice include: User Datagram Protocol Internet Protocol (UDP/IP) header compression, variable rate voice coding, and class-based queuing concepts. The DVI ATD Phase 2 effort demonstrated these advanced, integrated networking concepts over a constrained bandwidth wireless network to support simultaneous real-time and asynchronous datagram services.				
14. SUBJECT TERMS Integrated networking Resource reservation Packet voice Network compression			15. NUMBER OF PAGES 18	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

CONTENTS

INTRODUCTION	1
BACKGROUND	1
Network Voice Delivery Requirements	2
Design Approach	3
NETWORK VOICE APPLICATION	4
NRL Network Voice Application Background	4
Voice Algorithm Issues and Interface Development	4
DATA AND VOICE INTEGRATION ATD DIGITAL VOICE SUPPORT	6
Resource Reservation and Virtual Circuit Support	6
UDP/IP Header Compression	7
QoS-Aware API	9
IVOX ATD Subscriber Operation	11
Present and Future Design Efforts	12
SUMMARY	13
ACKNOWLEDGMENTS	14
REFERENCES	14

REAL-TIME PACKET VOICE TRAFFIC SUPPORT FOR INTEGRATED MOBILE MILITARY INTERNETWORKING

INTRODUCTION

The Naval Research Laboratory (NRL) has developed technology to support integrated multimedia services in military mobile communication networks. This report describes the NRL design and development of support for integrated real-time services within a self-organizing, mobile packet data network. The use of real-time service reservation to support integrated, low data rate voice communications is discussed. These efforts were achieved under the Office of Naval Research (ONR) Data and Voice Integration Advanced Technology Demonstration (DVI ATD) project. The primary project objective was to demonstrate the replacement of single, dedicated mobile communication networks with an open systems networking approach to support multiple uses such as tactical data broadcast, file transfer, image transfer, e-mail, command and control, group planning applications, and real-time interactive voice. The operation and design of various related software subsystems to support the real-time interactive voice feature are described in this document. Particularly, this paper emphasizes how the network resource reservation mechanism in the DVI ATD project supports voice and other real-time service integration over bandwidth-constrained wireless packet networks.

BACKGROUND

Traditional data network environments, such as those based upon the Defense Advanced Research Projects Agency (DARPA) Internet Protocol (IP) suite, provide many benefits to end users; however these networks typically provide only datagram "best effort" delivery service or a very simple quality-of-service (QoS) capability [1]. Circuit-switched communication protocols and networks generally provide service guarantees but lack in providing efficient service to bursty sources, adapting to link failures, and supporting "one-to-many" multipoint communications. There has been considerable recent research in developing an integrated services network architecture for traditional IP-based datagram networks [2]. This new capability will provide better support for applications, such as video, voice, multimedia conferencing, scientific visualization, and distributed computing within future network architectures.

The recent development of numerous high data rate network voice applications for use over traditional IP-based networks attests to a growing demand and interest in integrating real-time communications over existing data networks. The present use of high data rate packet voice digitization within network voice applications places unnecessary demands on performance and limits participation for mobile users and network sites with relatively low bandwidth resources. Typically, only "best effort" datagram delivery service is provided through present commercial IP routing products. NRL has conceived the notion of coupling low rate data network voice communications with a packet network traffic reservation mechanism as an enabling technology for bandwidth-constrained network users. Real-time voice calls or conferences can operate over this integrated network with guaranteed performance while maintaining minimum impact on other asynchronous data applications. Here we examine issues involved in developing and supporting integrated voice

communications within the DVI ATD project. In addition, we will discuss the development and use of a novel network resource reservation technique.

Network Voice Delivery Requirements

Previous digital voice communication systems have integrated robust error handling within the speech coding algorithm because the voice terminal application data have always had a direct interface to the link/physical communication layer. In contrast, the application layer within connectionless datagram networks maintains independence from the underlying communication media. In a global internetwork, this independence allows applications to communicate peer-to-peer across multiple, heterogeneous media. Lower layer protocols at the transport, network, and link layers usually assume the major responsibility for any error handling. Datagrams in error are either automatically retransmitted or dropped depending upon the protocol set used. As a result, the application layer (i.e., the voice terminal, in our case) does not usually incur bit errors in its communication data stream but may need to provide measures for handling undelivered packets and reconstruct packet ordering.

Internetwork Architecture

Many current digital voice compression algorithms have been targeted for voice communications over dedicated, synchronous communication channels. New design issues and performance factors arise when applying voice compression algorithms to packet communication networks. As more systems migrate to IP-compatible networks, contention for shared resources, especially external communications, will increase. Normally, processing of network data traffic is accomplished on a first-in, first-out basis. Applying this approach to overloaded external communications will be a significant step backwards in terms of service for users who are accustomed to dedicated circuits. For example, a full duplex 9600 baud dedicated circuit may seem more responsive than a 128 kbps circuit with 500 kbytes of data queued for transmission. To a user who requires low latency, the dedicated circuit sounds better. Unfortunately, dedicated circuits waste bandwidth while not in use, and most flagships are being configured with larger bandwidth circuits shared across multiple applications. A mechanism is required to accommodate applications requiring low latency and guaranteed throughput.

Integrated Resource Management Model

With underlying connectionless datagram delivery, there can be a significant amount of nondeterministic variance in the interpacket arrival times of a data stream. Furthermore, it is possible that data packets are received in a different order from which they were transmitted [3]. Current IP service provides best-effort delivery where all data flow obtains a similar QoS. Research is being conducted and initial test systems are in place for IP routing and data delivery techniques that bound factors affecting the quality of service of specific data flows, such as delay variance [4]. Such vendor-independent resource management techniques will provide proactive internetwork bandwidth allocation among and between application data streams.

The DVI ATD has successfully demonstrated an early capability of an integrated resource reservation technique to support real-time data streams over low bandwidth subnetworks within an IP-compliant architecture [5]. The combination of virtual circuit and datagram delivery models is used at the subnetwork to integrate simultaneous support of both guaranteed resource reservation and best-effort delivery services [6]. An application programming interface (API) is provided as mechanism for negotiating resource reservation at the application layer. Low data rate network voice communication has been the primary test application for this service, yet the model is easily extended to other integrated, real-time network traffic sources (e.g., common awareness data).

Design Approach

In the following section, we describe the design approach used to achieve a flexible network voice application that supports a number of ongoing research projects. We shall specifically discuss the design features supporting the a real-time network voice capability demonstrated within the DVI ATD effort.

Voice Digitization Requirements

The emphasis of the DVI ATD is to demonstrate an enhanced technical capability that can be applied to existing narrowband tactical networks to effectively integrate data and voice networking technologies. Targeted tactical data networks typically consist of relatively low bandwidth radio data links ranging from 2.4 kbps-1 Mbps rates. The DVI ATD Phase 1 and Phase 2 experiments have been successfully completed. These demonstrations targeted a tactical network consisting of 2.4 kbps data links as a baseline over which to support networked data/voice integration techniques. The subsystem techniques developed scale well to higher data rate networks. Supporting integrated voice communication over these low bandwidth subnetworks necessitates the use of advanced voice compression techniques to achieve voice digitization throughput rates less than 2400 bps. We chose to build upon the Federal Standard (FS) 1015 2400 bps LPC-10 algorithm [7] and a set of 600, 800, 1200 bps vector quantized LPC-based vocoders [8] for a prototype network voice traffic demonstration. The intelligible quality of these vocoders decreases with data rate, but such lower rate operation was necessary to demonstrate the integration concepts over 2.4 kbps communication links. The network voice application and support services are easily adaptable to any vocoder rate, so higher rate vocoder rates should be applied where more network bandwidth is available.

Software-Only Operation

We choose to target a software-only solution to provide existing commercial off-the-shelf (COTS) workstations with low data rate network voice capability. A software-only solution equates to maximum flexibility, and this work has been targeted to support numerous ongoing network research projects. System audio device drivers and analog interfaces were used to minimize the system support requirements on most workstations [5].

Efficient Network Use

The major design goal had been to achieve very low data rate network voice communication via advanced voice digitization and efficient use of network communication protocols. To achieve the real-time delivery requirements of digital voice communications, available reliable transport protocols (e.g., Transmission Control Protocol (TCP)) have been discarded in favor of a simple, unreliable transport mechanism (i.e., User Datagram Protocol (UDP)). When applied to the DVI ATD subsystem, the additional enhancement of UDP/IP header compression and the mapping of a UDP/IP delivery mechanism to a subnetwork virtual circuit mechanism resulted in efficient, robust end system to end system delivery.

Variable Rate Capability

In the past, digital voice communication systems have provided essentially continuous synchronous delivery of voice data across a dedicated communication channel. The dedicated communication channel is typically designed to provide a fixed amount of bandwidth, and vocoding algorithms have been designed to provide the best voice quality for bit rates supported within this bandwidth. In the connectionless networking environment where communication bandwidth is dynamically shared in an asynchronous fashion among distributed users and applications, adaptive rate voice coding techniques (e.g., silence detection) can improve bandwidth use dramatically. Adaptive rate voice coding can provide the same end-to-end voice quality as constant bit rate schemes, while maintaining a lower average network throughput. The synchronous nature of many

current voice encoding schemes has led to a belief that voice communication requires "stream" oriented data communications when, in fact, the information content of conversational voice is bursty in nature. NRL has designed an adaptive rate enhancement to existing Department of Defense (DoD) voice digitization algorithms and uses this as the default mode for network voice communication [9].

Within the DVI ATD, the variable voice capability has provided improved bandwidth use over the subnetworks due to the loose resource reservation mechanism applied. For a constant bit rate (CBR) service, the ATD subsystem will provide guaranteed bandwidth reservation, while with a variable rate service, unused bandwidth (e.g., voice silence periods, muted microphone) is made available to other datagram services.

NETWORK VOICE APPLICATION

NRL Network Voice Application Background

To meet some of the network voice design goals described earlier, NRL's Communication Systems Branch has created a low data rate Interactive Voice Exchange (IVOX) application that provides interactive, real-time internetwork voice communication on a number of standard computer workstations [10]. IVOX currently supports linear predictive coding (LPC)-based vocoders operating at speech data rates of 600-2400 bps. The 2400 bps LPC algorithm is a Federal Standard 1015 implementation, and the lower data rates (e.g., 600, 800, 1200 bps) are obtained via vector quantization of this basic analysis routine [11]. Within each vocoder algorithm, IVOX provides adaptive silence detection to further reduce the average data rate of each algorithm. For example, with adaptive silence detection, the 2400 bps LPC algorithm was shown to generate approximately 1100 bps of average throughput traffic during a typical conversational speech segment.

IVOX communicates over the Internet using the User Datagram Protocol (UDP) encapsulated in Internet Protocol (IP) packets. The UDP/IP suite provides unreliable, unordered best-effort delivery of data packets. An in-band signaling protocol is used to negotiate call setup and for subsequent communication session control. IVOX also provides a number of user controllable parameters for evaluating low data rate voice communication with connectionless datagram delivery. Example parameters include:

- number of vocoder frames per packet;
- resequencing window time setting;
- half-duplex or full-duplex operation; and
- real-time and non-real time interactive modes.

IVOX was designed as a research tool for very low data rate global internetwork voice communications. It is presently being applied in U.S. and NATO military packet switched radio network research and is in operational use within some distributed naval networks (e.g., Joint Deployed Intelligence Support System). A modified version of IVOX has been used as the voice application within the Phase 2 DVI ATD demonstrations.

Voice Algorithm Issues and Interface Development

Use and Enhancement of the LPC-10E Algorithm

Figure 1 is a block diagram of the LPC-10 vocoder. Linear prediction analysis is performed at a 22.5-ms frame rate by an open-loop 10th-order covariance method. IVOX uses this algorithm directly for constant bit rate (CBR) service at 2400 bps. The author has provided a variable rate-processing enhancement to the LPC-10 algorithm described below.

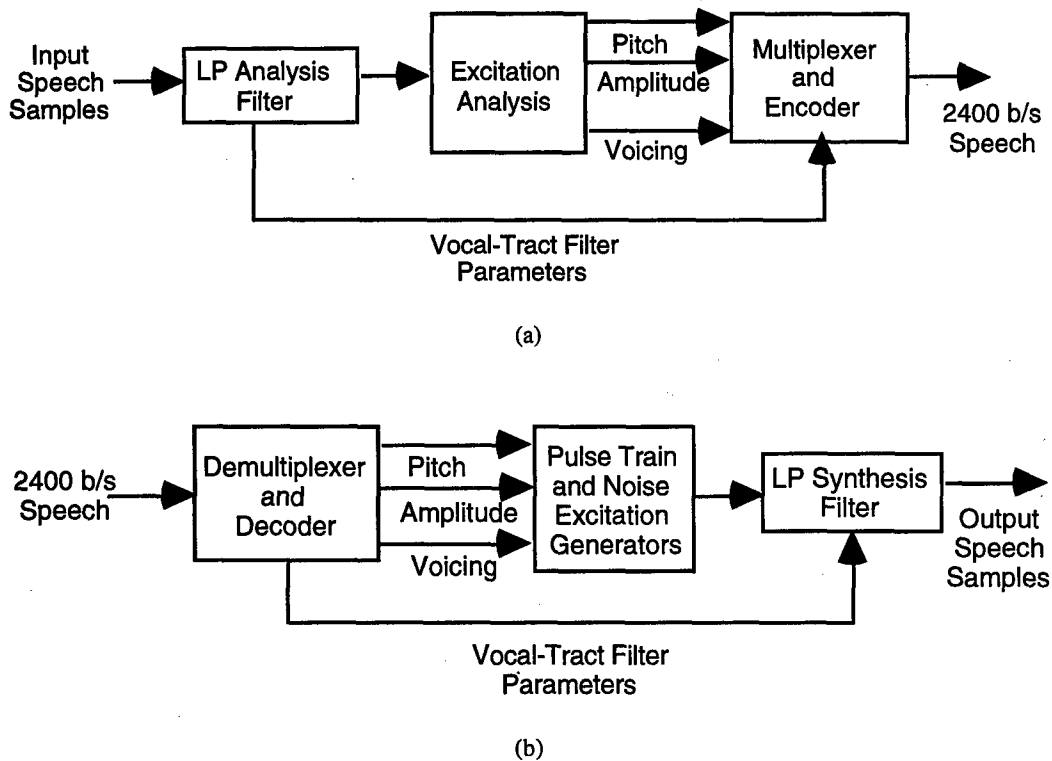


Fig. 1 - LPC-10 voice processor (a) transmitter and (b) receiver

There are two important frame-by-frame estimation parameters used to perform variable rate encoding decisions, energy, and voicing. A frame-by-frame root mean square (rms) energy measurement is already implemented by the LPC analysis routine, and its output is used to perform silence detection processing. In addition, unvoiced LPC frames contain redundant information that can be removed for network voice applications, resulting in a shorter frame length. The rationale and method will be discussed in the following paragraphs.

The FS 1015 LPC-10 standard employs a Hamming error correction coding method to improve performance in high bit error rate environments. As discussed earlier, within a network environment, the transport, network, and link layers generally provide delivery of error-free datagram data, although order and arrival times are not always guaranteed. Therefore, the Hamming error correction coding serves little purpose in the network voice application. We use this rationale to create a variable frame structure and allow silence frames to be dropped. If the energy value for an unvoiced frame is below a threshold value, it is considered a silence frame and need not be transmitted. The resulting vocoder transmission source will produce vocoder frames during "voice spurt" periods in which the frame rms energy is above a predefined threshold. To make this work, the Hamming decoder within LPC-10 has been deactivated, and silence frame processing has been added. Initial experiments and operational use with this variable rate vocoding scheme have demonstrated good intelligibility, as well as efficient use of the available transmission bandwidth. Over 1-s interval measurements, typical voice conversations have indicated an average transmission rate requirement of 1000-1200 b/s.

Vector Quantization Compression Methods

Work at the National Security Agency (NSA) was done to apply vector quantization methods to standard LPC-10 analysis to further compress voice data with a minimum increase in overall distortion [12]. Vector codebooks are being applied that result in overall constant bit rate throughputs of approximately 600, 800, and 1200 bps. Figure 2 describes the data flow.




Fig. 2 - Vector quantization processor

Fast tree search methods and hierarchical codebooks have been applied to vector quantization coding to significantly reduce the processing speed requirements over brute force single codebook approaches. This has made it possible to achieve software-only operation for very large effective codebook sizes (e.g., 25 bit) on standard workstations in real time. Memory requirements have also been reduced by the use of multiple parallel codebooks.

DATA AND VOICE INTEGRATION ATD DIGITAL VOICE SUPPORT

This section discusses the DVI ATD architecture and the mechanisms employed to support integrated packet voice.

Resource Reservation and Virtual Circuit Support

Within the DVI ATD subsystem, network integration of simultaneous data and voice is achieved by supporting classes of service at the subnetwork layer combined with a hybrid link layer forwarding mechanism (i.e., virtual circuit and datagram). Integrating a real-time application (e.g., voice) requiring guaranteed latency and bandwidth is accomplished through a resource reservation mechanism and a mapping of the data stream to virtual circuit link layer forwarding service.

Subnetwork Layer Cell-based Multiplexing and Queue Management

The DVI ATD has developed a flexible data packet multiplexing scheme to support integrated services over low bandwidth networks. Typical datagram packet sizes produced by end systems and applications are highly variable. Within the IP internetwork framework, typical interface maximum transmission unit (MTU) sizes are on the order of 1000 bytes or greater. These packet sizes are too large to support effective real-time multiplexing of data through low bandwidth networks (e.g., 2400 bps). At the network and data link layer interface, datagrams from each end system are segmented into 36-byte cells. This size was chosen as a reasonable compromise for limiting delay while achieving efficient message overhead ratios for low data rate links (e.g., 2400 bps). Virtual circuit and datagram cells are defined. Cells supporting virtual circuits can be distinguished from datagram cells by header content. There are two reasons to distinguish cell types. First, virtual circuit cells can be forwarded immediately at end points without reassembly, while datagram cells are buffered and reassembled prior to message delivery. Second, virtual circuit cells can be identified by the cell-forwarding queue process and are given special attention to ensure that reserved service quality is maintained (e.g., guaranteed bandwidth throughput).

The network voice application makes use of the subnetwork virtual circuit service to ensure integrated, real-time operation while coexisting with other random data application operation (e.g., e-mail, file transfer, collaborative planning).

Queue Management Strategy

The DVI ATD provides a queue management scheme for servicing the unique cell forwarding requirements of different traffic types or classes. A set of virtual circuit queues is managed in parallel with a datagram service queue that includes priority-based forwarding. Virtual circuit queues can be administered on a deterministic schedule to provide guaranteed latency service. A local table is maintained describing the flow requirements, negotiated at setup time, for each virtual circuit service entry. The datagram queue will be serviced after virtual circuit queues have received effective forwarding service. Within the datagram queue, packets can be further prioritized on an individual basis.

Link Layer Virtual Circuit Support

Within the DVI ATD high frequency (HF) subnetwork, the link layer operates using the NRL-developed Multichannel Architecture (MCA) protocol. MCA provides an interface to the virtual circuit setup protocol and performs subnetwork scheduling and virtual circuit management to maintain the underlying guaranteed service required by the voice application data streams [7].

UDP/IP Header Compression

In addition to the services discussed above, protocol compression has been necessary to support distributed network voice operation within the ATD demonstration. Relatively large headers are present in UDP, TCP, and IP protocols to meet the needs for a large address space in a global network, to support large number of communication services provided by workstations, and to support flexible functionality within these communication layers. Header compression or prediction techniques can be readily applied on limited resource links or subnets that will be supporting a relatively small number of simultaneous users and services. Applications that can most benefit from header prediction include transmit packets with a data payload size that is of the same order or small relative to network header size. Such applications include:

- telnet sessions;
- interactive voice and video;
- whiteboarding sessions;
- distributed interactive simulation;
- distributed computing; and
- tactical data (e.g. radar tracks or fire control).

Other applications such as file transfer or e-mail can gain performance on resource-limited links if relatively small MTU sizes are used within the link or subnetwork.

There are well-known techniques and standards in practice to compress TCP/IP datagram headers over serial and point-to-point links [13]. These methods are effective for *telnet* and *ftp* services based upon TCP transport methods, especially over bandwidth limited links (e.g., wireless, dial-up access). In the past, researchers paid little attention to developing standard prediction methods for UDP/IP datagrams, because they have felt the datagrams are found too infrequently and there was insufficient datagram-to-datagram coherence in place. This situation is changing with the proliferation of multicast services, resource reservation, and military-based real-time interactive applications using UDP/IP. The following emerging networking features further justify the need for standard UDP/IP header compression techniques: UDP stream applications are increasing in number (e.g., multicast, real-time voice/video, military tactical applications), many compressed data and interactive applications have significant header-to-data overhead ratios, and resource reservation methods are being designed for the Internet (e.g., Reservation Setup Protocol (RSVP)) and for limited military communication wireless links (e.g., DVI ATD virtual circuits).

The DVI ATD has developed and tested a UDP/IP compression method for use over virtual circuit connections. This method provides significant bandwidth improvements for compressed voice traffic (e.g., 800-2400 bps). This technique is based upon a priori information gained from the ATD resource reservation (i.e., virtual circuit) setup protocol, which is then used to reconstruct a UDP/IP datagram header at the receive node. Figure 3 shows a typical UDP/IP header and the resulting 6-byte compressed UDP/IP header used over a virtual circuit connection.

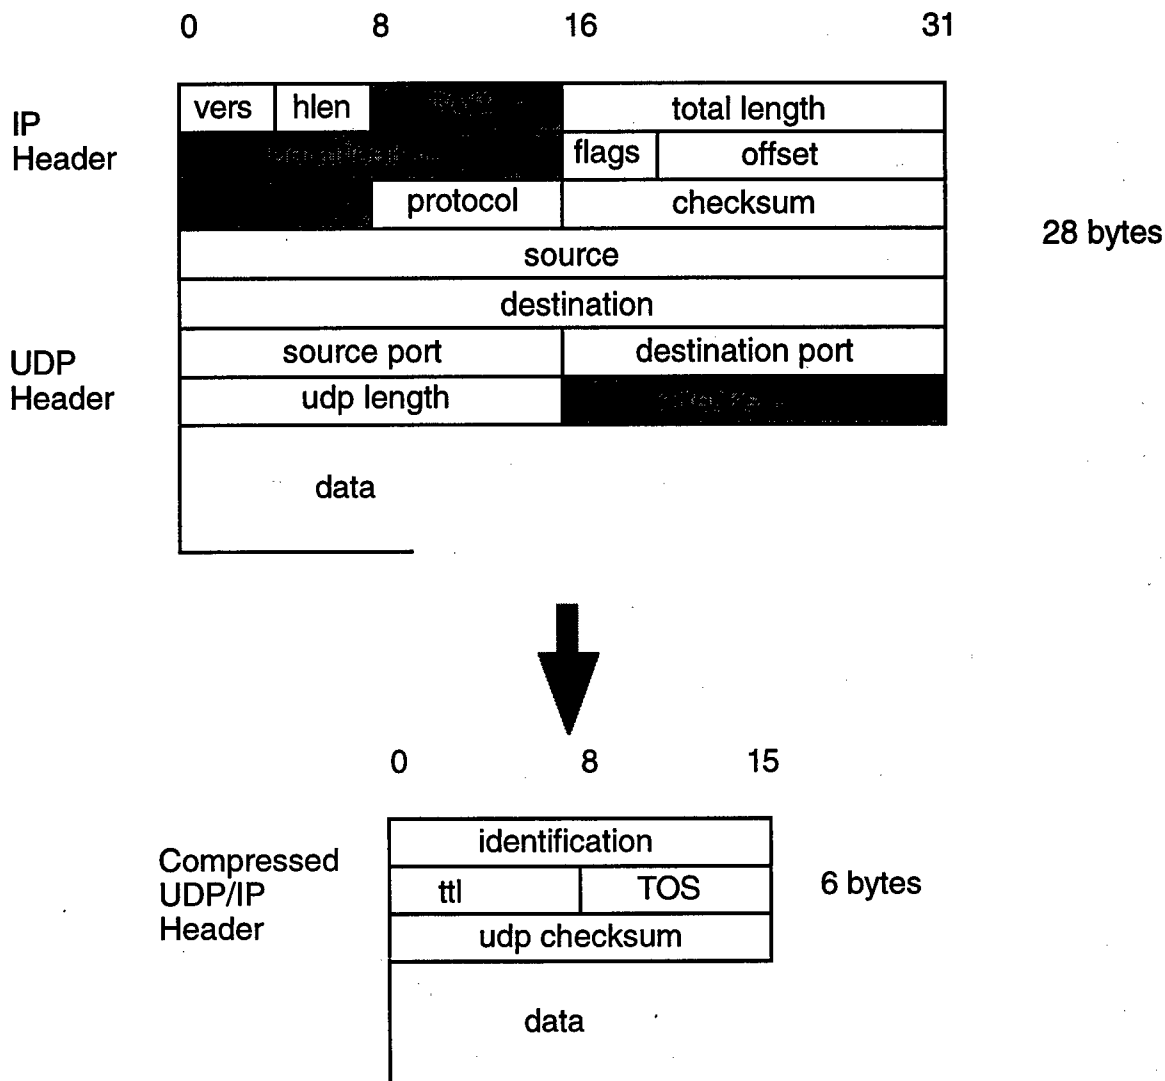


Fig. 3 - UDP/IP compression scheme

At the receive node, a UDP/IP header is reconstructed based on a priori information gained from knowledge of the assigned *service id* for a specific virtual circuit. Port numbers and IP source and destination addresses are mapped in from the service table information. Option and protocol id fields are considered to be constant. The IP checksum is recalculated as are the various length fields. Data integrity is preserved by transmitting the UDP checksum in the compressed header. The integrity of the compressed UDP/IP header is presently assumed to be sufficiently covered by the lower link layer protocol. Some additional modification may be required for increased operational robustness.

To demonstrate the utility of UDP/IP header compression over bandwidth-limited links, consider an example of a DVI ATD 800 bps voice application. The datagram payload can be as small as 18 bytes in order to minimize delay. The size of a standard UDP/IP message needed to carry this payload is 46 bytes. The bandwidth use in terms of network header to data payload ratio is 39%. With UDP/IP compression, the use is increased to 75%. Figure 4 is a comparison of bandwidth use over a range of data payload sizes. The region of typical compressed voice payload size shows significant improvement in bandwidth utilization (30% to 40 %).




Fig. 4 - Bandwidth use

Numerous typical military UDP/IP applications are anticipated to have data payloads around 10-100 bytes per message (e.g., compressed voice, interactive planning tools, tactical data traffic). Figure 4 demonstrates the usefulness of UDP/IP compression in these circumstances. This may be a critical enabling technical factor for a number of applications running in a military Internet with resource reservation.

QoS-Aware API

To support real-time voice and other applications requiring special services, the DVI ATD project developed host-side API software that is linked at compile time and used by a program as a resource management interface through socket and/or stream connections to an associated communication server device. Figure 5 shows the connections supporting an end system to communication server connection.

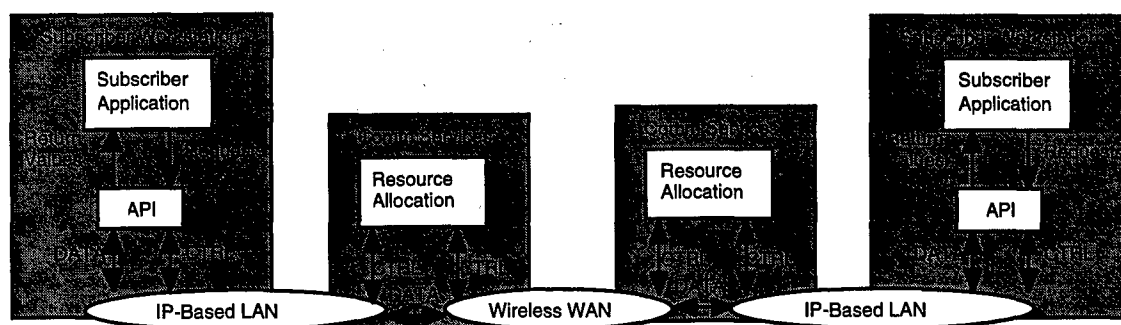


Fig. 5 - Data/Voice Integration Phase 2 API/CS architecture

The subscriber API is responsible for providing a data transfer, resource reservation setup, and resource access control interface. The API has been designed to run over TCP and UDP services, using socket or STREAMS/TLI facilities. There are a number of possible communication modes of operation using the API, and many of these modes were developed to map directly onto subnetwork layer services provided by MCA virtual circuit and datagram modes.

Functions

A number of basic API calls are provided to allow applications to interface and communicate with the ATD communication server subsystem. Table 1 lists API calls and provides a brief description of each call.

Table 1 - Subscriber API Calls

API Call	Description
csRegister	Registers new service with the communication server and provides basic service parameters.
csInfoRequest	Obtain registered service information.
csAccept	Accept connection requests from other subscribers.
csConnect	Connect to a peer subscriber offering a specific service.
csSend	Send data to peer subscriber.
csReceive	Receive data from peer subscriber.
csControl	Control virtual circuit parameters.
csSelect	Allows handling of multiple simultaneous connections.
csClose	Close a session.
csGetFd	Get the file descriptor for a specific session id.
csListen	Listen for service connection requests
csDeregister	Deregister a service offered by a subscriber.

IVOX ATD Subscriber Operation

This section discusses the particular operation of the NRL IVOX network voice application as applied to the DVI ATD network testbed. The IVOX network interface software has been successfully modified to make use of the DVI ATD subscriber API to allow resource reservation and virtual circuit support for real time voice calls. Each IVOX subscriber registers a service and then listens for connection requests from peer subscribers. Upon receiving a connection request, the IVOX server process will copy the reservation parameters and have its client process attempt a return connect to the calling subscriber. If the return connect is successful, communication can begin between peer subscribers. Figure 6 demonstrates the process.

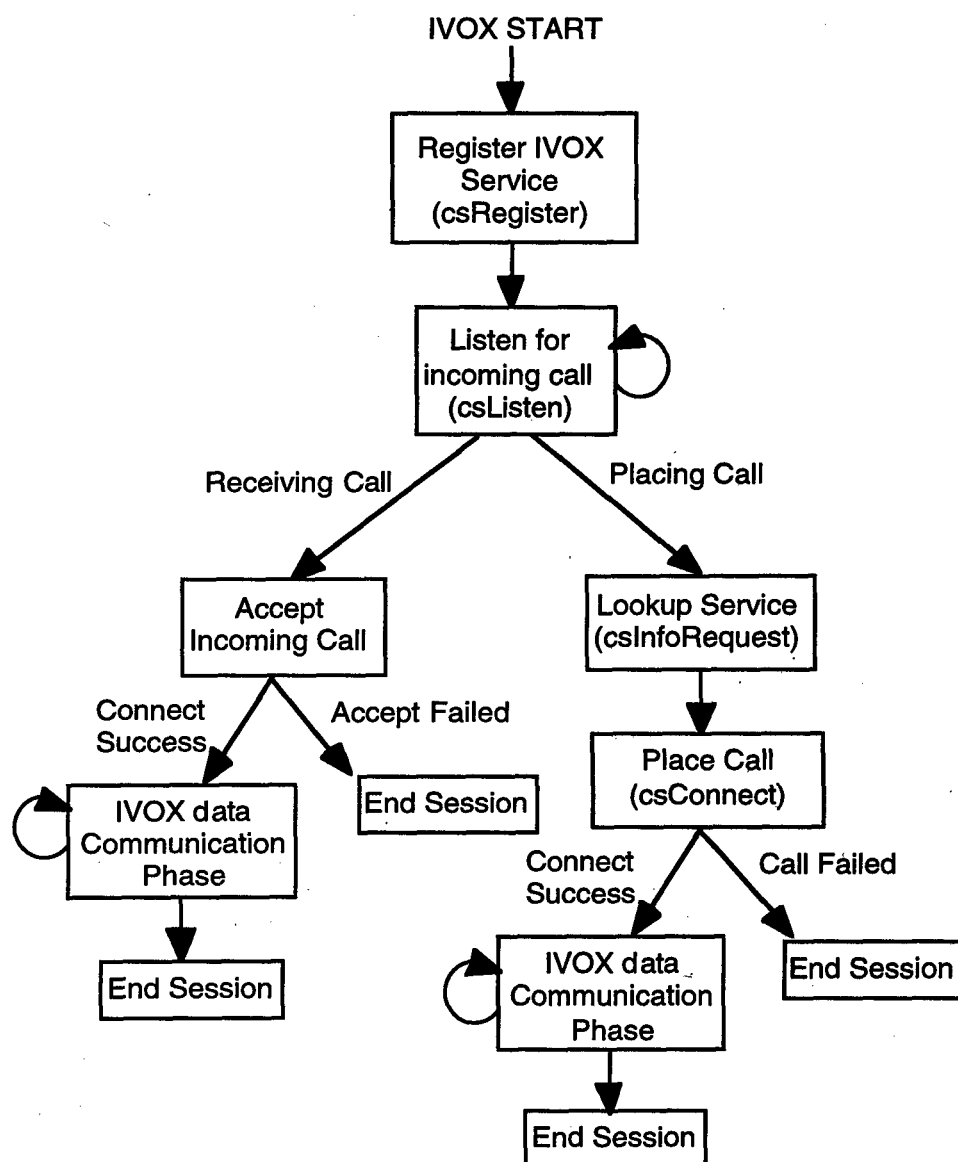


Fig. 6 - IVOX API state diagram

Resource Reservation Setup

Upon placing or servicing a connection request, the IVOX application will use the provided subscriber API to establish connection service parameters and negotiate for network resources. Within the ATD subsystem, the IVOX voice subscriber will ask that a virtual circuit service be

established to support the real-time delivery requirements. API virtual circuit requests are processed by the communication server and, if accepted, are passed on to the subnet access controller for subnetwork virtual circuit establishment. The IVOX voice subscriber uses the UDP transport service and is mapped into a subnetwork virtual circuit service that is either simplex, half duplex, or full duplex.

At present, the ATD resource reservation setup is source-based, with the source subscriber providing the reservation parameters for the connection to the network interface. In Phase 3, a receiver-based resource reservation scheme will be used based upon the emerging RSVP protocol.

Resource Reservation Enforcement

Upon obtaining a successful virtual circuit reservation, an end system IVOX voice exchange can take place. The quality of service parameters negotiated during the circuit setup phase are used to enforce traffic policing at forwarding nodes to ensure the voice application does not receive more than its fair share of bandwidth. A sliding window average is maintained of the effective throughput required to service a particular data stream. If the end application sources more data than allocated by a particular reservation, traffic policing will be enforced by dropping excess packets. If the traffic input rate is below the reserved amount for a service, this unused bandwidth becomes available for use by other asynchronous datagram services.

With this present resource management model, it is the responsibility of the voice application to not exceed allocated bandwidth during virtual circuit operation. Applying constant bit rate voice is relatively straightforward: the application asks for the bandwidth required and provides a constant source rate over time. Variable rate voice operation provides an opportunity for additional reservation tactics. The conservative approach is to request the maximum bandwidth required, which equates to that of constant bit rate operation and guarantees effective service. This can work efficiently with the DVI ATD reservation enforcement policy that will dynamically reallocate unused bandwidth to other services. The second option is to reserve an average bandwidth produced by human speech when using the variable rate algorithm. Earlier in this paper, we discussed the fact that this is usually significantly less than the maximum required bandwidth. Although this approach only provides service guarantee to a "well-behaved" variable rate application, it can be quite useful to provide a service that would otherwise be rejected by the admission control policy. We have used this approach in network experiments to support an LPC-10 (2400 bps) adaptive rate network voice stream, with only a 1500 bps reservation.

Present and Future Design Efforts

Phase 2 of the NRL DVI ATD project has been successful in demonstrating the feasibility of resource management techniques and protocols to support real-time services over limited bandwidth, wireless media. This technology demonstration has been accomplished through enhanced subnetwork protocols (i.e., MCA) and improved QoS data classification and forwarding methods based upon the IP protocol family (e.g., NRL/ViaSat communication server and API [14]).

Since the start of the DVI ATD project, the IETF and the associated research community have shown increased interest in this subject area and have recently demonstrated significant progress towards future standard protocol approaches. At the time of this writing, RSVP is being prepared for standards track submittal and there are several existing software implementations. Along with multicast support, RSVP implementation and interoperability are Phase 3 design goals for the DVI ATD project. RSVP support will provide numerous benefits: future interoperability, improved multicast support, scalability, and improved distributed operation. Given these benefits, there are still many moving targets within the emerging RSVP standard. Present RSVP API implementations work much in the same fashion as the ATD-developed API, but a key difference is that the RSVP API (rapi) to RSVP daemon (rsvpd) interface resides locally on the host machine. Therefore, data and control interface communication occurs internal to the host and can be better managed and

synchronized. The rsvpd translates API messages to/from network RSVP messages received over a standard datagram network. Applications do not generally see RSVP messages. Figure 7 shows the end system and rsvpd API architecture.

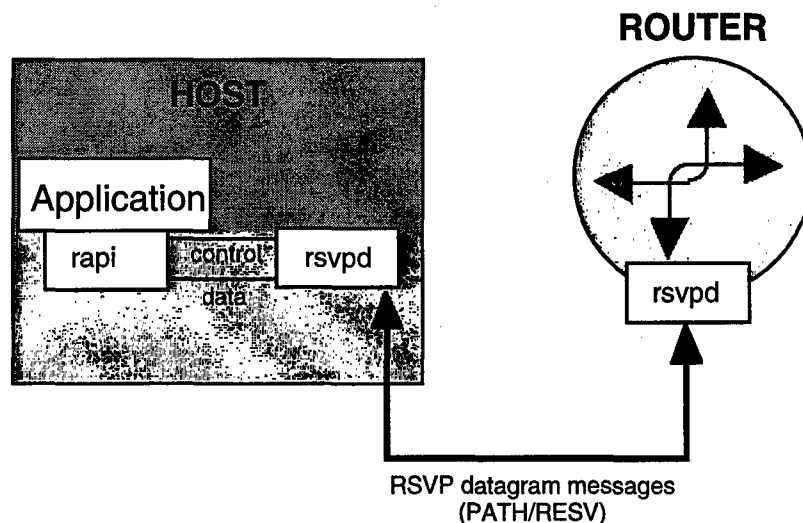


Fig. 7 - RAPI/RSVPD operation

Multicast Support

Multicasting is a technique used to forward copies of a single packet to a subset of all possible destinations within a network. When this subset is the entire set of all possible destinations, it is called broadcasting. The technique is presently supported on many LAN systems—most notable, Ethernet and Fiber Distributed Data Interface (FDDI). On these physical networks, a single multicast packet to all the members of a multicast group has the same overhead as a unicast packet to just one of them. This provides a significant advantage for many multimedia, distributed simulation, tactical, and collaborative applications by reducing the traffic bandwidth requirements dramatically. Multicasting technology has recently been extended to the wide-area network (WAN) environment through the introduction of multicast routing technology. IVOX presently contains support for standard IP multicast transport, including Internet Group Management Protocol (IGMP). A present Phase 3 ATD goal is to extend the Phase 2 architecture to support multicast routing. Further work will be required to integrate a multicast resource reservation mechanism (e.g., RSVP) and to develop an appropriate conference management scheme to support voice conferencing over low data rate tactical networks.

SUMMARY

The DVI ATD has combined an NRL-developed low data rate packet voice application (600-2400 bps) with a network resource reservation scheme to demonstrate the feasibility of an integrated tactical network architecture supporting real-time services. Effective real-time voice and data integration has been demonstrated over a subnetwork consisting of very low data rate links (approximately 2400 bps raw link rate) while maintaining a standard IP protocol end user network interface. This system was demonstrated throughout the summer of 1995 over a secure, HF wireless subnetwork within the Chesapeake Bay region. Absolute latency bounds and bandwidth reservation for network voice calls were preserved while simultaneously supporting nondeterministic, asynchronous datagram traffic routing (e.g., e-mail, ftp, collaborative whiteboarding).

With the increasing proliferation of distributed collaborative tools, situational awareness data, and integrated C4I networking for the warfighter, real-time network voice applications will provide a service of increasing value. Future efforts to enhance the ATD support for digital voice include the adaptation and implementation of an RSVP-based API interface, higher rate vocoding algorithms for

use in higher bandwidth networks, and multicasting capability. These modifications will improve the scalability of the DVI ATD networking architecture and align its network software components for interoperability with emerging COTS standards.

ACKNOWLEDGMENTS

The author thanks R. Brian Adamson for his integration of API enhancements to the NRL IVOX application to support the DVI ATD demonstrations. Tom Tremain, Mary Kohler, and John Collura of the NSA R224 deserve special thanks for their technical assistance and development of prototype vector quantized LPC algorithms and real-time implementations used by this project. My gratitude to the Office of Naval Research (ONR) ATD office and SPAWAR PD-70 for sponsoring and supporting this work. Finally, thanks go out to Ed Althouse and the entire NRL Data and Voice Integration ATD team for their hard work and contributions to make integrated, low data rate network voice support a reality against considerable technical risk.

REFERENCES

1. D. Clark, "The Design Philosophy of the DARPA Internet Protocols," Proceedings of the ACM SIGCOMM, August 1988.
2. L.Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala, "RSVP: A New Resource Reservation Protocol," *IEEE Networks Magazine*, September 1993.
3. W.R. Stevens, *TCP/IP Illustrated*, Vol. 1 (Addison-Wesley Publishing Company, Reading, Massachusetts, 1994).
4. S. Floyd, "Link-Sharing and Resource Management Models for Packet Networks," Technical Report, Lawrence Berkeley Laboratory, December 14, 1994.
5. E.L. Althouse, J.P. Macker, J.P. Hauser, and D.J. Baker, "Integrated Services in Tactical Communication Systems," *1995 NRL Review*, Naval Research Laboratory, Washington, D.C., pp. 141-143.
6. W. Thoet, D.J. Baker, and D.N. McGregor, "A Multichannel Architecture for Naval Task Force Communication," NRL/FR/5520--94-9703, Naval Research Laboratory, Washington, D.C., January 30, 1994.
7. Federal Standard 1015, "Analog to Digital Conversion of Voice by 2400 bits/s Linear Predictive Coding," published by the General Services Administration, November 28, 1984.
8. M.A. Kohler, "Real Time Low Rate Voice Coding," Version 1.0, Technical Report TR-R22-10-93, National Security Agency, Ft. Meade, Md., May 18, 1993.
9. J. Collura, *Modern Methods of Speech Processing* (Kluwer Academic Press, Norwell, MA, 1995), Chap. 2.
10. J. P. Macker and R.B. Adamson, "I VOX—The Interactive Voice eXchange Application," NRL FR 5520--96-9805, Naval Research Laboratory, Washington, D.C., Feb 20, 1996.
11. J.P. Macker and R.B. Adamson, "Linear Predictive Coding (LPC) Speech 'C' Software Library," NRL Technical Memorandum 5520-65A, Naval Research Laboratory, Washington, D.C., April 28, 1994.
12. J.P. Macker and R.B. Adamson, "A Variable Rate Voice Coder using LPC-10E," NRL Technical Memorandum 5520-92, Naval Research Laboratory, Washington, D.C., July 12, 1994.

13. V. Jacobson, "Compressing TCP/IP Headers for Low-speed Serial Links," Internet Request for Comments 1144, February 1990.
14. "Software Requirements Specification for the Data and Voice Integration Advanced Technology Demonstration", ViaSat Document #608300-94-001, ViaSat Inc., Carlsbad, CA., June 1994.